

UNITED STATES PATENT APPLICATION

FOR

**IMPROVING THE RECEIVE PERFORMANCE OF A NETWORK ADAPTER  
BY DYNAMICALLY TUNING ITS INTERRUPT DELAY**

Inventor(s): Daniel R. Gaur

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP  
12400 Wilshire Boulevard, 7<sup>th</sup> Floor  
Los Angeles, California 90025  
(425) 827-8600

"Express Mail" Label Number EL431685448US

Date of Deposit June 6, 2001

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. §1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

  
\_\_\_\_\_  
Sharon E. Farnus

IMPROVING THE RECEIVE PERFORMANCE OF A NETWORK ADAPTER  
BY DYNAMICALLY TUNING ITS INTERRUPT DELAY

5

TECHNICAL FIELD  
OF THE INVENTION

This disclosure relates to network adapters, and more particularly, but not exclusively, to apparatus and methods of improving the receive performance of a network adapter by dynamically tuning the network adapter's interrupt delay based on an evaluation of an incoming network load.

10  
15  
20  
25  
30  
35  
40  
45  
50  
55  
60  
65  
70  
75  
80  
85  
90  
95  
100  
105  
110  
115  
120  
125  
130  
135  
140  
145  
150  
155  
160  
165  
170  
175  
180  
185  
190  
195  
200  
205  
210  
215  
220  
225  
230  
235  
240  
245  
250  
255  
260  
265  
270  
275  
280  
285  
290  
295  
300  
305  
310  
315  
320  
325  
330  
335  
340  
345  
350  
355  
360  
365  
370  
375  
380  
385  
390  
395  
400  
405  
410  
415  
420  
425  
430  
435  
440  
445  
450  
455  
460  
465  
470  
475  
480  
485  
490  
495  
500  
505  
510  
515  
520  
525  
530  
535  
540  
545  
550  
555  
560  
565  
570  
575  
580  
585  
590  
595  
600  
605  
610  
615  
620  
625  
630  
635  
640  
645  
650  
655  
660  
665  
670  
675  
680  
685  
690  
695  
700  
705  
710  
715  
720  
725  
730  
735  
740  
745  
750  
755  
760  
765  
770  
775  
780  
785  
790  
795  
800  
805  
810  
815  
820  
825  
830  
835  
840  
845  
850  
855  
860  
865  
870  
875  
880  
885  
890  
895  
900  
905  
910  
915  
920  
925  
930  
935  
940  
945  
950  
955  
960  
965  
970  
975  
980  
985  
990  
995  
1000  
1005  
1010  
1015  
1020  
1025  
1030  
1035  
1040  
1045  
1050  
1055  
1060  
1065  
1070  
1075  
1080  
1085  
1090  
1095  
1100  
1105  
1110  
1115  
1120  
1125  
1130  
1135  
1140  
1145  
1150  
1155  
1160  
1165  
1170  
1175  
1180  
1185  
1190  
1195  
1200  
1205  
1210  
1215  
1220  
1225  
1230  
1235  
1240  
1245  
1250  
1255  
1260  
1265  
1270  
1275  
1280  
1285  
1290  
1295  
1300  
1305  
1310  
1315  
1320  
1325  
1330  
1335  
1340  
1345  
1350  
1355  
1360  
1365  
1370  
1375  
1380  
1385  
1390  
1395  
1400  
1405  
1410  
1415  
1420  
1425  
1430  
1435  
1440  
1445  
1450  
1455  
1460  
1465  
1470  
1475  
1480  
1485  
1490  
1495  
1500  
1505  
1510  
1515  
1520  
1525  
1530  
1535  
1540  
1545  
1550  
1555  
1560  
1565  
1570  
1575  
1580  
1585  
1590  
1595  
1600  
1605  
1610  
1615  
1620  
1625  
1630  
1635  
1640  
1645  
1650  
1655  
1660  
1665  
1670  
1675  
1680  
1685  
1690  
1695  
1700  
1705  
1710  
1715  
1720  
1725  
1730  
1735  
1740  
1745  
1750  
1755  
1760  
1765  
1770  
1775  
1780  
1785  
1790  
1795  
1800  
1805  
1810  
1815  
1820  
1825  
1830  
1835  
1840  
1845  
1850  
1855  
1860  
1865  
1870  
1875  
1880  
1885  
1890  
1895  
1900  
1905  
1910  
1915  
1920  
1925  
1930  
1935  
1940  
1945  
1950  
1955  
1960  
1965  
1970  
1975  
1980  
1985  
1990  
1995  
2000  
2005  
2010  
2015  
2020  
2025  
2030  
2035  
2040  
2045  
2050  
2055  
2060  
2065  
2070  
2075  
2080  
2085  
2090  
2095  
2100  
2105  
2110  
2115  
2120  
2125  
2130  
2135  
2140  
2145  
2150  
2155  
2160  
2165  
2170  
2175  
2180  
2185  
2190  
2195  
2200  
2205  
2210  
2215  
2220  
2225  
2230  
2235  
2240  
2245  
2250  
2255  
2260  
2265  
2270  
2275  
2280  
2285  
2290  
2295  
2300  
2305  
2310  
2315  
2320  
2325  
2330  
2335  
2340  
2345  
2350  
2355  
2360  
2365  
2370  
2375  
2380  
2385  
2390  
2395  
2400  
2405  
2410  
2415  
2420  
2425  
2430  
2435  
2440  
2445  
2450  
2455  
2460  
2465  
2470  
2475  
2480  
2485  
2490  
2495  
2500  
2505  
2510  
2515  
2520  
2525  
2530  
2535  
2540  
2545  
2550  
2555  
2560  
2565  
2570  
2575  
2580  
2585  
2590  
2595  
2600  
2605  
2610  
2615  
2620  
2625  
2630  
2635  
2640  
2645  
2650  
2655  
2660  
2665  
2670  
2675  
2680  
2685  
2690  
2695  
2700  
2705  
2710  
2715  
2720  
2725  
2730  
2735  
2740  
2745  
2750  
2755  
2760  
2765  
2770  
2775  
2780  
2785  
2790  
2795  
2800  
2805  
2810  
2815  
2820  
2825  
2830  
2835  
2840  
2845  
2850  
2855  
2860  
2865  
2870  
2875  
2880  
2885  
2890  
2895  
2900  
2905  
2910  
2915  
2920  
2925  
2930  
2935  
2940  
2945  
2950  
2955  
2960  
2965  
2970  
2975  
2980  
2985  
2990  
2995  
3000  
3005  
3010  
3015  
3020  
3025  
3030  
3035  
3040  
3045  
3050  
3055  
3060  
3065  
3070  
3075  
3080  
3085  
3090  
3095  
3100  
3105  
3110  
3115  
3120  
3125  
3130  
3135  
3140  
3145  
3150  
3155  
3160  
3165  
3170  
3175  
3180  
3185  
3190  
3195  
3200  
3205  
3210  
3215  
3220  
3225  
3230  
3235  
3240  
3245  
3250  
3255  
3260  
3265  
3270  
3275  
3280  
3285  
3290  
3295  
3300  
3305  
3310  
3315  
3320  
3325  
3330  
3335  
3340  
3345  
3350  
3355  
3360  
3365  
3370  
3375  
3380  
3385  
3390  
3395  
3400  
3405  
3410  
3415  
3420  
3425  
3430  
3435  
3440  
3445  
3450  
3455  
3460  
3465  
3470  
3475  
3480  
3485  
3490  
3495  
3500  
3505  
3510  
3515  
3520  
3525  
3530  
3535  
3540  
3545  
3550  
3555  
3560  
3565  
3570  
3575  
3580  
3585  
3590  
3595  
3600  
3605  
3610  
3615  
3620  
3625  
3630  
3635  
3640  
3645  
3650  
3655  
3660  
3665  
3670  
3675  
3680  
3685  
3690  
3695  
3700  
3705  
3710  
3715  
3720  
3725  
3730  
3735  
3740  
3745  
3750  
3755  
3760  
3765  
3770  
3775  
3780  
3785  
3790  
3795  
3800  
3805  
3810  
3815  
3820  
3825  
3830  
3835  
3840  
3845  
3850  
3855  
3860  
3865  
3870  
3875  
3880  
3885  
3890  
3895  
3900  
3905  
3910  
3915  
3920  
3925  
3930  
3935  
3940  
3945  
3950  
3955  
3960  
3965  
3970  
3975  
3980  
3985  
3990  
3995  
4000  
4005  
4010  
4015  
4020  
4025  
4030  
4035  
4040  
4045  
4050  
4055  
4060  
4065  
4070  
4075  
4080  
4085  
4090  
4095  
4100  
4105  
4110  
4115  
4120  
4125  
4130  
4135  
4140  
4145  
4150  
4155  
4160  
4165  
4170  
4175  
4180  
4185  
4190  
4195  
4200  
4205  
4210  
4215  
4220  
4225  
4230  
4235  
4240  
4245  
4250  
4255  
4260  
4265  
4270  
4275  
4280  
4285  
4290  
4295  
4300  
4305  
4310  
4315  
4320  
4325  
4330  
4335  
4340  
4345  
4350  
4355  
4360  
4365  
4370  
4375  
4380  
4385  
4390  
4395  
4400  
4405  
4410  
4415  
4420  
4425  
4430  
4435  
4440  
4445  
4450  
4455  
4460  
4465  
4470  
4475  
4480  
4485  
4490  
4495  
4500  
4505  
4510  
4515  
4520  
4525  
4530  
4535  
4540  
4545  
4550  
4555  
4560  
4565  
4570  
4575  
4580  
4585  
4590  
4595  
4600  
4605  
4610  
4615  
4620  
4625  
4630  
4635  
4640  
4645  
4650  
4655  
4660  
4665  
4670  
4675  
4680  
4685  
4690  
4695  
4700  
4705  
4710  
4715  
4720  
4725  
4730  
4735  
4740  
4745  
4750  
4755  
4760  
4765  
4770  
4775  
4780  
4785  
4790  
4795  
4800  
4805  
4810  
4815  
4820  
4825  
4830  
4835  
4840  
4845  
4850  
4855  
4860  
4865  
4870  
4875  
4880  
4885  
4890  
4895  
4900  
4905  
4910  
4915  
4920  
4925  
4930  
4935  
4940  
4945  
4950  
4955  
4960  
4965  
4970  
4975  
4980  
4985  
4990  
4995  
5000  
5005  
5010  
5015  
5020  
5025  
5030  
5035  
5040  
5045  
5050  
5055  
5060  
5065  
5070  
5075  
5080  
5085  
5090  
5095  
5100  
5105  
5110  
5115  
5120  
5125  
5130  
5135  
5140  
5145  
5150  
5155  
5160  
5165  
5170  
5175  
5180  
5185  
5190  
5195  
5200  
5205  
5210  
5215  
5220  
5225  
5230  
5235  
5240  
5245  
5250  
5255  
5260  
5265  
5270  
5275  
5280  
5285  
5290  
5295  
5300  
5305  
5310  
5315  
5320  
5325  
5330  
5335  
5340  
5345  
5350  
5355  
5360  
5365  
5370  
5375  
5380  
5385  
5390  
5395  
5400  
5405  
5410  
5415  
5420  
5425  
5430  
5435  
5440  
5445  
5450  
5455  
5460  
5465  
5470  
5475  
5480  
5485  
5490  
5495  
5500  
5505  
5510  
5515  
5520  
5525  
5530  
5535  
5540  
5545  
5550  
5555  
5560  
5565  
5570  
5575  
5580  
5585  
5590  
5595  
5600  
5605  
5610  
5615  
5620  
5625  
5630  
5635  
5640  
5645  
5650  
5655  
5660  
5665  
5670  
5675  
5680  
5685  
5690  
5695  
5700  
5705  
5710  
5715  
5720  
5725  
5730  
5735  
5740  
5745  
5750  
5755  
5760  
5765  
5770  
5775  
5780  
5785  
5790  
5795  
5800  
5805  
5810  
5815  
5820  
5825  
5830  
5835  
5840  
5845  
5850  
5855  
5860  
5865  
5870  
5875  
5880  
5885  
5890  
5895  
5900  
5905  
5910  
5915  
5920  
5925  
5930  
5935  
5940  
5945  
5950  
5955  
5960  
5965  
5970  
5975  
5980  
5985  
5990  
5995  
6000  
6005  
6010  
6015  
6020  
6025  
6030  
6035  
6040  
6045  
6050  
6055  
6060  
6065  
6070  
6075  
6080  
6085  
6090  
6095  
6100  
6105  
6110  
6115  
6120  
6125  
6130  
6135  
6140  
6145  
6150  
6155  
6160  
6165  
6170  
6175  
6180  
6185  
6190  
6195  
6200  
6205  
6210  
6215  
6220  
6225  
6230  
6235  
6240  
6245  
6250  
6255  
6260  
6265  
6270  
6275  
6280  
6285  
6290  
6295  
6300  
6305  
6310  
6315  
6320  
6325  
6330  
6335  
6340  
6345  
6350  
6355  
6360  
6365  
6370  
6375  
6380  
6385  
6390  
6395  
6400  
6405  
6410  
6415  
6420  
6425  
6430  
6435  
6440  
6445  
6450  
6455  
6460  
6465  
6470  
6475  
6480  
6485  
6490  
6495  
6500  
6505  
6510  
6515  
6520  
6525  
6530  
6535  
6540  
6545  
6550  
6555  
6560  
6565  
6570  
6575  
6580  
6585  
6590  
6595  
6600  
6605  
6610  
6615  
6620  
6625  
6630  
6635  
6640  
6645  
6650  
6655  
6660  
6665  
6670  
6675  
6680  
6685  
6690  
6695  
6700  
6705  
6710  
6715  
6720  
6725  
6730  
6735  
6740  
6745  
6750  
6755  
6760  
6765  
6770  
6775  
6780  
6785  
6790  
6795  
6800  
6805  
6810  
6815  
6820  
6825  
6830  
6835  
6840  
6845  
6850  
6855  
6860  
6865  
6870  
6875  
6880  
6885  
6890  
6895  
6900  
6905  
6910  
6915  
6920  
6925  
6930  
6935  
6940  
6945  
6950  
6955  
6960  
6965  
6970  
6975  
6980  
6985  
6990  
6995  
7000  
7005  
7010  
7015  
7020  
7025  
7030  
7035  
7040  
7045  
7050  
7055  
7060  
7065  
7070  
7075  
7080  
7085  
7090  
7095  
7100  
7105  
7110  
7115  
7120  
7125  
7130  
7135  
7140  
7145  
7150  
7155  
7160  
7165  
7170  
7175  
7180  
7185  
7190  
7195  
7200  
7205  
7210  
7215  
7220  
7225  
7230  
7235  
7240  
7245  
7250  
7255  
7260  
7265  
7270  
7275  
7280  
7285  
7290  
7295  
7300  
7305  
7310  
7315  
7320  
7325  
7330  
7335  
7340  
7345  
7350  
7355  
7360  
7365  
7370  
7375  
7380  
7385  
7390  
7395  
7400  
7405  
7410  
7415  
7420  
7425  
7430  
7435  
7440  
7445  
7450  
7455  
7460  
7465  
7470  
7475  
7480  
7485  
7490  
7495  
7500  
7505  
7510  
7515  
7520  
7525  
7530  
7535  
7540  
7545  
7550  
7555  
7560  
7565  
7570  
7575  
7580  
7585  
7590  
7595  
7600  
7605  
7610  
7615  
7620  
7625  
7630  
7635  
7640  
7645  
7650  
7655  
7660  
7665  
7670  
7675  
7680  
7685  
7690  
7695  
7700  
7705  
7710  
7715  
7720  
7725  
7730  
7735  
7740  
7745  
7750  
7755  
7760  
7765  
7770  
7775  
7780  
7785  
7790  
7795  
7800  
7805  
7810  
7815  
7820  
7825  
7830  
7835  
7840  
7845  
7850  
7855  
7860  
7865  
7870  
7875  
7880  
7885  
7890  
7895  
7900  
7905  
7910  
7915  
7920  
7925  
7930  
7935  
7940  
7945  
7950  
7955  
7960  
7965  
7970  
7975  
7980  
7985  
7990  
7995  
8000  
8005  
8010  
8015  
8020  
8025  
8030  
8035  
8040  
8045  
8050  
8055  
8060  
8065  
8070  
8075  
8080  
8085  
8090  
8095  
8100  
8105  
8110  
8115  
8120  
8125  
8130  
8135  
8140  
8145  
8150  
8155  
8160  
8165  
8170  
8175  
8180  
8185  
8190  
8195  
8200  
8205  
8210  
8215  
8220  
8225  
8230  
8235  
8240  
8245  
8250  
8255  
8260  
8265  
8270  
8275  
8280  
8285  
8290  
8295  
8300  
8305  
8310  
8315  
8320  
8325  
8330  
8335  
8340  
8345  
8350  
8355  
8360  
8365  
8370  
8375  
8380  
8385  
8390  
8395  
8400  
8405  
8410  
8415  
8420  
8425  
8430  
8435  
8440  
8445  
8450  
8455  
8460  
8465  
8470  
8475  
8480  
8485  
8490  
8495  
8500  
8505  
8510  
8515  
8520  
8525  
8530  
8535  
8540  
8545  
8550  
8555  
8560  
8565  
8570  
8575  
8580  
8585  
8590  
8595  
8600  
8605  
8610  
8615  
8620  
8625  
8630  
8635  
8640  
8645  
8650  
8655  
8660  
8665  
8670  
8675  
8680  
8685  
8690  
8695  
8700  
8705  
8710  
8715  
8720  
8725  
8730  
8735  
8740  
8745  
8750  
8755  
8760  
8765  
8770  
8775  
8780  
8785  
8790  
8795  
8800  
8805  
8810  
8815  
8820  
8825  
8830  
8835  
8840  
8845  
8850  
8855  
8860  
8865  
8870  
8875  
8880  
8885  
8890  
8895  
8900  
8905  
8910  
8915  
8920  
8925  
8930  
8935  
8940  
8945  
8950  
8955  
8960  
8965  
8970  
8975  
8980  
8985  
8990  
8995  
9000  
9005  
9010  
9015  
9020  
9025  
9030  
9035  
9040  
9045  
9050  
9055  
9060  
9065  
9070  
9075  
9080  
9085  
9090  
9095  
9100  
9105  
9110  
9115  
9120  
9125  
9130  
9135  
9140  
9145  
9150  
9155  
9160  
9165  
9170  
9175  
9180  
9185  
9190  
9195  
9200  
9205  
9210  
9215  
9220  
9225  
9230  
9235  
9240  
9245  
9250  
9255  
9260  
9265  
9270  
9275  
9280  
9285  
9290  
9295  
9300  
9305  
9310  
9315  
9320  
9325  
9330  
9335  
9340  
9345  
9350  
9355  
9360  
9365  
9370  
9375  
9380  
9385  
9390  
9395  
9400  
9405  
9410  
94

which may have arrived during a scheduling latency created as the processor completes its current tasks and switches contexts to execute the interrupt handler. The interrupt handler comprises a special routine that is executed when a specific interrupt occurs, and which includes instructions for dealing with the particular situation that caused the interrupt. The interrupt 5 handler examines the network controller to determine the cause of the interrupt, for instance, the receipt of new frames from the network, and performs the necessary post-interrupt cleanup, which may include routing the incoming new frames to an appropriate application.

Each interrupt, and the execution of the interrupt handler, introduces an amount of “overhead” to the computer system’s processor. The overhead comprises work or information that provides support for a computing process, but is not an intrinsic part of the operation or data. For example, with each incoming packet, the processor may need to send a signal, comprising the packet, over a bus, or otherwise transfer the packet to one or more other components of the computer system. This overhead may significantly impact processing time, especially in the context of modern operating systems, such as Windows NT®, or Windows® 2000.

15 Two distinct scenarios are presently utilized by network adapters, via the network controller, to signal the successful receipt of a new frame: (1) the network controller may generate an interrupt immediately following the receipt of a new packet; or (2) the network controller may delay the generation of an interrupt following the receipt of a new packet. The delay may be a set period of time, or it may be defined by the receipt of a pre-defined number of 20 packets, or a fixed timeout period. By delaying generation of the interrupt, the network controller has an opportunity to receive additional frames before interrupting, thereby effectively “bundling” together several packets into a single interrupt. As a consequence of this “bundling,” the overhead associated with the interrupt is effectively amortized across several frames, thereby

lessening the overhead to the computer system's processor. However, an additional consequence of incorporating a delay prior to generation of the interrupt, is that the average latency per packet is increased. Latency refers to the period of time that passes between a point in time at which a packet arrives at the network adapter from the network, and a point in time at which the interrupt 5 handler is executed.

While delayed interrupts exhibit better system performance under heavy incoming network loads because multiple packets can be handled with the overhead of a single interrupt, immediate interrupts exhibit better system performance under light incoming network loads because each packet incurs a minimal latency, and the associated overhead of additional 10 interrupts is not significant enough to degrade system performance.

Current devices, including network adapters, utilize a fixed policy for determining when to generate an interrupt. As mentioned previously, a network adapter may generate an interrupt, via the network controller, either immediately upon receiving a new frame, or the network adapter may introduce a delay before generating the interrupt. In either case, the 15 interrupt policy is static, meaning that once the policy is initialized, the network adapter will continue to interrupt at a defined, fixed rate regardless of network or system conditions.

While a fixed interrupt scheme will always work, in the sense that the network adapter will carry out its intended functions, such a fixed interrupt scheme cannot account for variations in the incoming network traffic density, or in the computer system. As such, a static 20 interrupt policy may be optimized for a particular ideal workload, but as runtime workload shifts away from the ideal workload, the static policy may degrade overall system performance, thereby causing excessive processor utilization (e.g., if the network adapter interrupts too often), or poor response latency (e.g., if the network adapter delays the interrupts too long).

**BRIEF DESCRIPTION OF THE  
VARIOUS VIEWS OF THE DRAWINGS**

In the drawings, like reference numerals refer to like parts throughout the various views of the non-limiting and non-exhaustive embodiments of the present invention, and

5 wherein:

Figure 1 is an illustration of an event timeline showing the occurrence of events in an example network environment in which incoming network traffic is relatively light;

Figure 2 is an illustration of an event timeline showing the occurrence of events in an example network environment in which incoming network traffic is relatively heavy;

Figure 3 is a flow diagram illustrating the implementation of an embodiment of the present invention; and

Figure 4 is a graphical representation illustrating how the interrupt delay may vary in an example network environment as incoming network traffic fluctuates.

DETAILED DESCRIPTION OF  
THE ILLUSTRATED EMBODIMENTS

Embodiments of a system and method for dynamically modifying the interrupt delay of a network adapter to optimize receive performance of the network adapter based on incoming network traffic are described in detail herein. In the following description, numerous specific details are provided, such as the identification of various system components, to provide a thorough understanding of embodiments of the invention. One skilled in the art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, etc. In still other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of various embodiments of the invention.

Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearance of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

As an overview, embodiments of the invention provide systems and methods for dynamically tuning the interrupt delay of a network adapter in response to an evaluation of incoming network traffic. In an embodiment, the network adapter, which may include a network controller, comprises a component of a computer system that may also include a processor and a device driver. The network adapter may be capable of being interconnected with the processor,

and further capable of being connected to a network from which incoming network traffic may be generated.

In various embodiments, a scheduling algorithm adjusts duration of the interrupt delay based on, for example, the number of packets received in conjunction with a previous 5 interrupt, or other suitable monitoring technique. As the duration of the interrupt delay is adjusted in response to incoming network traffic, the scheduling algorithm may also vary a pair of threshold values, corresponding to the upper and lower limits of the incoming network traffic load, for example, the number of packets per interrupt. These threshold values may then be utilized as a range for a future analysis of network traffic density, and a determination of when 10 and how the interrupt delay should be adjusted to optimize performance and system efficiency. Other features of the illustrated embodiments will be apparent to the reader from the foregoing and the appended claims, and as the detailed description and discussion is read in conjunction with the accompanying drawings.

Referring now to the drawings, and in particular to Figure 1, there is illustrated an 15 event timeline generally at 10 showing the occurrence of example events in an example network environment in which incoming network traffic is relatively light. The reader will appreciate that the term “relatively” as used throughout this disclosure with reference to network traffic loads as being relatively light or relatively heavy is not intended to define any specific network traffic pattern or density, but is intended only to illustrate embodiments of the present invention. 20 Network traffic that may be defined as relatively light in one instance, may be defined as relatively heavy in another.

In order to optimize efficiency in the computer system, and receive performance of the network adapter in a network environment, such as that illustrated in Figure 1, the network

controller will ideally interrupt after almost every packet received. By doing so, each incoming packet incurs only minor latencies, and the interrupts are too infrequent to substantially impact the performance of the processor, which may comprise a component of the computer system.

For example, a first packet (i) arrives at a point in time on the event timeline, indicated by

5 reference numeral 12. If, as illustrated in Figure 1, the interrupt delay is set at zero, the network controller will generate an interrupt immediately upon receiving the first packet (i) (at reference numeral 12). Following the interrupt is a “scheduling latency,” indicated by reference numeral 14. The scheduling latency 14 corresponds to a length of time it takes the computer system’s processor (*e.g.*, microprocessor or digital signal processor) to complete its current tasks and switch contexts to execute the interrupt handler. The scheduling latency 14 may vary with the tasks being undertaken by the processor at the time the interrupt is generated.

In the example event timeline illustrated in Figure 1, a second packet (i +1) arrives at a point in time on the event timeline, indicated by reference numeral 16, during the scheduling latency 14. At the point in time when the interrupt handler is executed (indicated by

15 reference numeral 18), both the first packet (i) and the second packet (i +1) are “cleaned up” by the network adapter’s device driver. The device driver comprises a software component that permits a computer system to communicate with a device, such as a network adapter, and which includes an interrupt handler for manipulating and handling incoming packets. Under network traffic conditions such as those illustrated in Figure 1, an interrupt delay would typically only 20 increase the packet latency, without any associated benefit in regard to bundling a number of packets together to be handled in a single interrupt.

Referring now primarily to Figure 2, there is illustrated an event timeline generally at 20 showing the occurrence of example events in an example network environment in

which incoming network traffic density is relatively heavy. In order to optimize efficiency in the computer system, and the receive performance of the network adapter under these network traffic conditions, the network controller will ideally interrupt after some defined period of delay, thereby allowing the computer system's processor and the device driver to handle a greater number of packets with the overhead of a single interrupt, while at the same time, only minimally affecting the latency of each incoming packet.

In the illustrated event timeline 20 of Figure 2, a first packet (j) arrives at the network adapter at a point in time indicated by reference numeral 22. The arrival of this first packet (j) triggers the start of the interrupt delay, indicated by reference numeral 24, that precedes generation of the interrupt by the network controller (occurring at the point in time indicated by reference numeral 30). As illustrated in Figure 2, two additional packets, (j + 1) and (j + 2), arrive at points in time along the event timeline 20, indicated by reference numerals 26 and 28 respectively. At the end of the interrupt delay (indicated by reference numeral 30), the network controller generates the interrupt, and the scheduling latency, indicated by reference numeral 32, begins in a manner similar to that described in conjunction with Figure 1. As mentioned previously, the scheduling latency will not necessarily be identical for each corresponding interrupt, but may vary with the current tasks being undertaken by the computer system's processor at the time the interrupt is generated by the network controller.

During the scheduling latency 32, three additional packets,  $(j + 3)$ ,  $(j + 4)$ , and  $(j + 5)$ , arrive at points in time along the event timeline 20, indicated by reference numerals 34, 36, and 38 respectively. At the point in time when the interrupt handler is executed (indicated by reference numeral 40), packets  $(j)$  through  $(j + 5)$  are “cleaned up” by the network adapter’s device driver. As the reader will appreciate from a comparison of Figures 1 and 2, assuming the

scheduling latency to be identical in both situations, without the interrupt delay provided in Figure 2, the network controller would have effectively missed packets (j + 4) and (j + 5) (indicated by reference numerals 36 and 38 respectively), and therefore would have had to interrupt the processor a second time to schedule and execute an interrupt handler to facilitate 5 management of these last two packets.

In optimizing an interrupt schedule for the network adapter, competing interests, between interrupting too often, thereby leading to excessive processor utilization, and not interrupting often enough, thereby leading to poor response latency, may be considered. As the reader will appreciate from the foregoing discussion in conjunction with Figures 1 and 2, a virtually infinite range of network traffic loads may exist, and systems and methods that utilize static interrupt scheduling policies are unable to adequately cope with changing workloads. 10

Having observed the example network environments of Figures 1 and 2, and the distinct methods for efficiently handling the arrival of incoming network traffic under different network traffic loads, attention may now be given to systems and methods of the illustrated 15 embodiments that facilitate the optimization of the network adapter's interrupt schedule by dynamically tuning the interrupt delay that precedes generation of the interrupt to signal the receipt of incoming network traffic. By adjusting the duration of the interrupt delay when, and as needed, the computer system of which the network adapter may be a component, may more efficiently handle incoming network traffic without excessive processor utilization, and without 20 poor response latency.

Turning our attention now primarily to Figure 3, an interrupt scheduling process, illustrating the implementation of a scheduling algorithm in accordance with an embodiment of the invention, is depicted as a flow diagram generally at 42. The scheduling algorithm may be

implemented as part of the device driver, or as a separate routine, executable by a controller, such as a microprocessor, a digital signal processor, or other device known to those skilled in the art. The reader will appreciate that the scheduling algorithm may be embodied as a set of instructions included on a machine-readable medium. A machine-readable medium includes any mechanism that provides (*e.g.*, stores and/or transmits) information in a form readable by a machine (*e.g.*, a computer). For example, a machine-readable medium includes read only memory (“ROM”); random access memory (“RAM”); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (*e.g.*, carrier tones, infrared signals, and digital signals); and the like.

The embodiment illustrated by the process 42 generally comprises a method by which a network controller’s interrupt delay is adjusted in response to a rate at which new packets are arriving from a network. The process 42 begins by monitoring the incoming network traffic load (see, *e.g.*, reference numeral 44) to generate a monitoring input for use as a comparison tool in the remainder of the process 42. The monitoring input comprises a value corresponding to the incoming network traffic load. Monitoring the incoming network traffic load may be accomplished, in one embodiment, via the utilization of statistical counters that periodically examine the network controller to ascertain the rate of incoming packets. In other embodiments, the incoming network traffic load may be determined via a calculated number of packets being received per interrupt, or by other suitable systems and methods known to those skilled in the art.

Following monitoring of the incoming network traffic load (see, *e.g.*, block 44), the monitoring input may be communicated to the device driver or other routine comprising the scheduling algorithm, and the process 42 proceeds to compare (see, *e.g.*, block 46) the incoming

network traffic load with an upper threshold. The upper threshold may be set at a value, in the first instance, by a system administrator based on an anticipation of incoming network traffic, or the upper threshold may be set at a predefined value upon system initialization. For example, the upper threshold may be set at a value between 64 and 512 packets per interrupt. It will be understood that this value may vary with the configuration of the network controller, the computer system with which the network controller may be implemented, and the incoming network traffic. As such, the values and ranges given herein are for illustrative purposes only, and should not be construed to limit the scope of the invention.

The reader will appreciate that the units (e.g., packets per interrupt), which define the value of the upper threshold, as well as a value corresponding to a lower threshold that will be discussed in greater detail hereinbelow, will correspond to the monitored incoming network traffic load. For example, if the incoming network traffic load is being monitored by the number of packets per interrupt, then the values of the upper and lower thresholds will also be defined by a number of packets per interrupt.

15 If the incoming network traffic load is greater than the value of the upper threshold, as determined by the comparison at block 46, then the process 42 proceeds to increase the network controller's interrupt delay (see, e.g., reference numeral 48). By increasing the interrupt delay, a greater number of incoming packets may be bundled together for processing via a single execution of the device driver's interrupt handler (see, e.g., Figure 2), thereby 20 enabling the computer system to spend more time processing other tasks, rather than repeatedly scheduling and servicing interrupt handlers.

Although the increase in the interrupt delay correspondingly increases the average latency per packet, the processor of the computer system may be able to more efficiently handle

the rising workload without unduly burdening the system. In an embodiment of the invention, the interrupt delay may vary between 0 and 128 ms, and each increase in the interrupt delay may correspond to an increase of from 3 to 5 ms, for example. The interrupt delay may be set, in a first instance, by a system administrator based on an anticipation of incoming network traffic, or 5 may be set at a predefined value upon system initialization. The setting for the interrupt delay may be correlated to the initial settings for the upper and lower thresholds as well to create a range of network traffic loads corresponding to the duration of the interrupt delay.

After increasing the interrupt delay (see, e.g., block 48), the process 42 proceeds to increase the value of the upper threshold (see, e.g., reference numeral 50). As mentioned previously, the upper threshold may be set, in an embodiment, at a value within the range of 64 to 512 packets per interrupt, for example. An increase in the upper threshold may correspond to an increase of from 28 to 36 packets per interrupt, for example. Following the increase in the value of the upper threshold (see, e.g., block 50), the value of the lower threshold is also increased (see, e.g., reference numeral 52) by an amount that may be equal to the amount of the increase in the value of the upper threshold, for example. In other embodiments, the upper and lower thresholds may be increased by different amounts, and in still other embodiments, the upper threshold may be increased while the lower threshold remains unchanged. By increasing the value of both the upper threshold and the lower threshold by an equal amount, a defined range is maintained corresponding to a particular interrupt delay. When the rate of incoming network traffic falls outside of this range, then the interrupt delay may again be adjusted accordingly, and the upper and lower thresholds may once again be adjusted. Following adjustment of the value of the upper and lower thresholds, the process 42 ends, awaiting the next monitoring input.

Where the incoming network traffic load is less than or equal to the upper threshold, as determined by the comparison at block 46, the process 42 proceeds to compare (see, *e.g.*, reference numeral 54) the incoming network traffic load with the lower threshold. As with the upper threshold discussed above, the lower threshold corresponds to a value that may vary, in 5 an embodiment, between 64 and 512 packets per interrupt, for example, and may be set, in the first instance, by a system administrator or upon system initialization. Again, the value of the lower threshold will depend on the configuration of the network controller, the computer system with which the network controller may be implemented, and the incoming network traffic load. As such, the range of values set forth above should not be construed to limit the scope of the 10 present invention.

If the incoming network traffic load is greater than or equal to the lower threshold (as determined by the comparison at block 54), then the process 42 ends, awaiting the next monitoring input. If the incoming network traffic load is less than the lower threshold (as determined by the comparison at block 54), then the process 42 proceeds to decrease the network 15 controller's interrupt delay (see, *e.g.*, reference numeral 56). By decreasing the interrupt delay (see, *e.g.*, block 56), the computer system's processor and the device driver can handle the received packets more quickly, thereby decreasing the latency associated with each packet. The interrupt rate may be increased by this decrease in the interrupt delay, but the frequency of the interrupts may not be so great as to degrade system performance. In an embodiment of the 20 invention, the interrupt delay may vary between 0 and 128 ms, while the decrease in the interrupt delay may correspond, in an embodiment, to a decrease of from 1 to 3 ms, for example.

After decreasing the interrupt delay at block 56, the process 42 proceeds to decrease the value of the lower threshold (see, *e.g.*, reference numeral 58). As mentioned

previously, the value of the lower threshold may, in an embodiment, be set within the range of 64 to 512 packets per interrupt, for example. The decrease in the value of the lower threshold (see, e.g., block 58) may, in an embodiment, correspond to a decrease of from 4 to 12 packets per interrupt, for example. Following the decrease in the value of the lower threshold at block 58,

5 the value of the upper threshold may also be decreased (see, e.g., reference numeral 60) by an amount that may be equal to the amount of the decrease in the value of the lower threshold, for example. In other embodiments, the upper and lower thresholds may be decreased by different amounts, and in still other embodiments, the lower threshold may be decreased while the upper threshold remains unchanged. As with the increases in the values of the upper and lower

10 thresholds discussed previously, by decreasing the value of both the lower threshold and the upper threshold by an equal amount, a defined range is maintained corresponding to a particular interrupt delay. When the rate of incoming network traffic falls outside of this range, then the interrupt delay may again be adjusted accordingly, and the lower and upper thresholds may once again be adjusted. Following adjustment of the lower and upper threshold values (at blocks 58

15 and 60 respectively), the process 42 ends, awaiting the next monitoring input.

Referring now primarily to Figure 4, a graphical representation illustrating how the interrupt delay may vary in an example network environment is shown generally at 62. In the example illustrated in Figure 4, the interrupt delay (represented by the line having reference numeral 63) may be set at 60 ms, for example, at the point in time indicated by reference numeral 64. At this same point in time 64, the upper threshold may be set at 300 packets per interrupt, for example, and the lower threshold may be set at 268 packets per interrupt, for example. These values may be preset, as discussed previously in regard to Figure 3, by a system administrator or upon system initialization, or may represent any point in time at which the

interrupt delay, and the upper and lower thresholds have reached the given values in response to incoming network traffic conditions.

As time passes, a second point in time, indicated by reference numeral 66, is reached, wherein the incoming network traffic load in the example network environment has 5 increased beyond the value of the upper threshold to, for example, 310 packets per interrupt. Because the incoming network traffic load has increased beyond the value of the upper threshold, the scheduling algorithm, an embodiment of which is described hereinabove with reference to Figure 3, increases the interrupt delay by 4 ms, for example, to 64 ms, and increases the values of the upper and lower thresholds by, for example, 32 packets per interrupt, to 332 10 packets per interrupt and 300 packets per interrupt, respectively.

Continuing on the time line to the point in time indicated by reference numeral 68, the incoming network traffic load in the example network environment has again increased beyond the value of the upper threshold to, for example, 353 packets per interrupt. As with the previous response, the scheduling algorithm increases the interrupt delay by 4 ms, for example, 15 to 68 ms, and increases the values of the upper and lower thresholds by, for example, 32 packets per interrupt, to 364 packets per interrupt and 332 packets per interrupt, respectively. Further continuing on the time line to the point in time indicated by reference numeral 70, the incoming network traffic load in the example network environment has decreased below the value of the lower threshold to, for example, 327 packets per interrupt. In response to this decreased 20 workload, the scheduling algorithm decreases the interrupt delay by 2 ms, for example, to 66 ms, and decreases the values of the upper and lower thresholds by, for example, 8 packets per interrupt, to 356 packets per interrupt and 324 packets per interrupt. This process of adjustment may continue as the incoming network traffic load varies, thereby optimizing the receive

performance of the network adapter and the efficiency of the computer system with which the network adapter may be interconnected.

It is possible that the incoming network traffic load may change so dramatically that the load, defined as packets per interrupt, for example, may remain above or below the value 5 of the upper or lower threshold, respectively, even after an adjustment is made by the scheduling algorithm. For instance, in the example given above in regard to the point in time designated by reference numeral 70, had the incoming network traffic load decreased to 321 packets per interrupt, for example, the adjustment in the value of the lower threshold to 324 packets per interrupt would not have encompassed the monitored network traffic load. In such 10 circumstances, assuming the level remains at a point below the lower threshold, then the scheduling algorithm will again adjust the interrupt delay and the values of the upper and lower thresholds in accordance with the previous examples. It will be appreciated that the example 15 values given in the preceding discussion and figures are for illustrative purposes only, and should not be construed to limit the scope of the present invention beyond that which is expressly set forth in the claims.

While the invention is described and illustrated here in the context of a limited number of embodiments, the invention may be embodied in many forms without departing from the spirit of the essential characteristics of the invention. The illustrated and described embodiments, including what is described in the abstract of the disclosure, are therefore to be considered in all respects as illustrative and not restrictive. The scope of the invention is indicated by the appended claims rather than by the foregoing description, and all changes which come within the meaning and range of equivalency of the claims are intended to be embraced therein.